RECEIVED
CENTRAL FAX CENTER
MAY 0 2 2006

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Glasco

Attorney Docket No.: NWISP030

Application No.: 10/608,846

Examiner: THAI, TUAN V.

Filed: June 27, 2003

Group: 2186

Title: METHODS AND APPARATUS FOR SENDING TARGETED PROBES

CERTIFICATE OF FACSIMILE TRANSMISSION
I hereby certify that this correspondence is being transmitted by facsimile to fax number 571-273-8300 to the U.S. Patent and Trademark Office on May 2, 2006

Signed: _____
Joyce P. Pereira

PRE-APPEAL BRIEF REQUEST FOR REVIEW

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Applicant requests review of the final rejection in the above-identified application. No amendments are being filed with this request.

This request is being filed with a Notice of Appeal.

The review is requested for the reasons stated on the attached sheets.

Remarks begin on page 6 of this paper.

IN THE CLAIMS

1. (Original) A computer system, comprising:

a home cluster including a first plurality of processing nodes and a home cache coherence controller, the first plurality of processing nodes and the home cache coherence controller interconnected in a point-to-point architecture;

a remote cluster including a second plurality of processing nodes and a remote cache coherence controller, the remote cache coherence controller configured to receive a probe from the home cluster, identify a processing node from the second plurality of processing nodes that owns a cache line corresponding to the probe, and send a targeted probe to the processing node.

2. (Original) The computer system of claim 1, wherein the processing node has the cache line in the owned or modified state.

3. (Original) The computer system of claim 1, wherein information for identifying the processing node that owns the cache line is provided in the probe from the home cluster.

4. (Original) The computer system of claim 1, wherein information for identifying the processing node that owns the cache line is provided by a coherence directory associated with the home cluster.

5. (Original) The computer system of claim 4, wherein the coherence directory maintains information on which clusters and processing nodes own particular cache lines.

6. (Original) The computer system of claim 1, wherein the remote cache coherence controller is further configured to send a directed probe to the processor that owns the cache line associated with the probe.

7. (Original) The computer system of claim 6, wherein the remote cache coherence controller is associated with a pending buffer.

8. (Original) The computer system of claim 7, wherein the remote cache coherence controller is set to receive a single response corresponding to the probe by setting the pending buffer.

9. (Original) The computer system of claim 8, wherein the probe is a read probe.

10. (Original) The computer system of claim 1, wherein the remote cache coherence controller does not send a directed probe if the cache line is also cached shared in the owning cluster.

11. (Original) The computer system of claim 1, further comprising a request cluster that generates a probe request triggering the probe from the home cluster

12. (Original) The computer system of claim 1, wherein each processing node comprises a processor, a memory controller, and a cache.

13. (Original) The computer system of claim 12, wherein each processing node has a portion of the computer system address space.

14. (Original) The computer system of claim 1, wherein the home cache coherence controller forwards the probe before probing home cluster processing nodes.

15. (Original) The computer system of claim 1, wherein the home cache coherence controller forwards the probe after sending probes to home cluster processing nodes.

16. (Previously Presented) A method for providing owning node information, the method comprising:

receiving a request for ownership of a memory line from a request cluster, the request cluster comprising a plurality of request cluster processing nodes and a request cache coherence controller interconnected in a point-to-point architecture;

identifying owning node information associated with the request for ownership at a home cluster, the home cluster comprising a plurality of home cluster processing nodes; and

maintaining owning node information in a coherence directory associated with the home cluster.

17. (Original) The method of claim 16, wherein the request for ownership of the memory line is a read block modify request.

18. (Original) The method of claim 16, wherein the request for ownership of the memory line is a change to dirty request.

19. (Original) The method of claim 16, wherein the request for ownership of the memory line is a validate block request.

20. (Original) The method of claim 16, further comprising maintaining owning cluster information in the coherence directory.

21. (Original) The method of claim 16, further comprising receiving a subsequent probe request from the request cluster.

22. (Original) The method of claim 16, further comprising determining if the state of a memory line associated with the subsequent probe is in the owned or modified state.

23. (Original) The method of claim 16, further comprising sending a targeted probe to an owning cluster if the state is owned or modified.

24. (Original) The method of claim 23, wherein the targeted probe includes owning node information.

25. (Original) The method of claim 24, wherein the targeted probe allows probing of a single processing node in the owning cluster.

26. (Previously Presented) An apparatus for providing owning node information, the apparatus comprising:

means for receiving a request for ownership of a memory line from a request cluster, the request cluster comprising a plurality of request cluster processing nodes and a request cache coherence controller interconnected in a point-to-point architecture;

means for identifying owning node information associated with the request for ownership at a home cluster, the home cluster comprising a plurality of home cluster processing nodes; and

means for maintaining owning node information associated with the home cluster.

27. (Original) The apparatus of claim 26, further comprising means for maintaining owning cluster information.

28. (Original) The apparatus of claim 26, further comprising means for receiving a subsequent probe request from the request cluster.

29. (Original) The apparatus of claim 26, further comprising means for determining if the state of a memory line associated with the subsequent probe is in the owned or modified state.

30. (Original) The apparatus of claim 26, further comprising means for sending a targeted probe to an owning cluster if the state is owned or modified.

31. (Original) The apparatus of claim 30, wherein the targeted probe includes owning node information.

32. (Original) The apparatus of claim 31, wherein the targeted probe allows probing of a single processing node in the owning cluster.

33. (Previously Presented) A computer readable medium comprising computer code for managing owning node information, the computer readable medium comprising:

computer code for receiving a request for ownership of a memory line from a request cluster, the request cluster comprising a plurality of request cluster processing nodes and a request cache coherence controller interconnected in a point-to-point architecture;

computer code for identifying owning node information associated with the request for ownership at a home cluster, the home cluster comprising a plurality of home cluster processing nodes; and

computer code for maintaining owning node information associated with the home cluster.

34. (Original) The apparatus of claim 33, further comprising computer code for maintaining owning cluster information.

35. (Original) The apparatus of claim 33, further comprising computer code for receiving a subsequent probe request from the request cluster.

REMARKS

Claims 1-35 are pending. Claims 1-35 were rejected. Claims 1-32 were rejected under 35 U.S.C. 102(e) as being anticipated by Edirisooriya (2003/0195939A1), hereinafter Edi. Claims 33-35 were rejected under 35 U.S.C. 103(a) as being unpatentable over Edi.

Edi describes a multiprocessor system that "includes a plurality of processors 12 and 14 that are communicatively coupled via an interconnection network 16. The processors 12 and 14 are implemented using any desired processing unit ... The interconnection network 16 is implemented using any suitable shared bus or other communication network or interface that permits multiple processors to communicate with each other and, if desired, with other system agents such as, for example, memory controllers ... Additionally, the processors 12 and 14 respectively include caches 22 and 24, cache controllers 26 and 28 and request queues 30 and 32." (Paragraphs 12-14) Edi also describes use of the MSI, MESI and MOESI protocols "to eliminate unnecessary data transfers between processors." (Paragraph 33) Edi does not mention multiple processor clusters and only describes multiple processors connected over an interconnection network such as a shared bus.

The Examiner argues that Edi states "persons of ordinary skill in the art will recognize that the multiprocessor system 10 may include additional processors or agents that are also communicatively coupled via the interconnection network 16, if desired." (Paragraph 18) Edi mentions that "while the interconnection network 16 is preferably implemented using a hardwired communication medium, other communication media, including wireless media, could be used instead." (Paragraph 13) The only example of an interconnection network Edi provides is a "shared bus." (Paragraphs 3 and 13)

Having multiple processors connected over an interconnection network does not teach or suggest clusters of processors. The independent claims 1, 16, 26 and 33 all recite a "home cluster" including a "plurality of processing nodes" and a "remote cluster" including a "plurality of processing nodes." A home cluster and a remote cluster are described throughout the present application and examples are depicted throughout in the Figures, e.g. Figures 1A, 1B, 2, 11, and 12 and associated description. Edi does not teach or suggest any home cluster or remote cluster. Furthermore, Edi does not teach or suggest any home cluster including a plurality of processing nodes or any remote cluster including a plurality of processing nodes. The Examiner argues that

a remote cluster is "other processing nodes 14" shown in Figure 1 of Edi. The Applicants respectfully disagree. Edi merely depicts a conventional architecture in Figure 1 and associated description with a processor 12 connected to a processor 14 over a shared bus 16. No clusters are shown. No clusters of processing nodes are shown. Even if we were to hypothetically add numerous additional processors to the Edi system, nothing in Edi teaches or suggests separating the processors into processor clusters. The flows charts in Figures 2 and 3 and the associated description merely depict interactions between multiple processors connected over a bus 16. Figures 4a-4d and associated description also similarly only show "various states through which the multiprocessor system 10 shown in FIG. 1 progresses." No clusters are depicted, taught, or even suggested.

Independent claim 1 also recites "the first plurality of processing nodes and the home cache coherence controller interconnected in a point-to-point architecture." Even though independent claims 16, 26, and 33 are believed patentable in their current form, independent claims 16, 26, and 33 have been amended to recite a plurality of request cluster processing nodes "and a request cache coherence controller interconnected in a point-to-point architecture." Edi only describes processors connected using a shared bus or interconnection network. A shared bus or interconnection network does not teach or suggest processing nodes and a cache coherence controller interconnected in a point-to-point architecture.

Furthermore, point-to-point links and clusters are not obvious modifications of a conventional architecture such as the one described in Edi. "By using point-to-point links instead of a conventional shared bus or external network, multiple processors are used efficiently in a system sharing the same memory space. Processing and network efficiency are also improved by avoiding many of the bandwidth and latency limitations of conventional bus and external network based multiprocessor architectures. According to various embodiments, however, linearly increasing the number of processors in a point-to-point architecture leads to an exponential increase in the number of links used to connect the multiple processors. In order to reduce the number of links used and to further modularize a multiprocessor system using a point-to-point architecture, multiple clusters are used." (page 6, lines 24-32)

In addition, claim 1 recites "identify a processing node from the second plurality of processing nodes that owns a cache line corresponding to the probe, and send a targeted probe to the processing node" and claims 16, 26, and 33 recite "identifying owning node information...

[and] maintaining owning node information associated with the home cluster." The Examiner argues that Edi describes determining whether the cache block associated with a request is in an owned state. However, Edi does not teach or suggest identifying the processing node that owns the cache line. Conventional systems such as Edi only provides MEOSI state information for a particular cache line and do not identify the node or processing node owning the cache line.

The techniques of the present invention recognize that this is beneficial for a variety of reasons. For example, "a coherence directory is used to eliminate the transmission of a request to a memory controller in a home cluster. A coherence directory can also be used to more accurately send targeted probes. In one example, only a node owning a particular memory line needs to be probed. Information can be added to probe requests and probes to identify the owning node and allow probes to be directed only at owning nodes in a given cluster." (page 9, lines 1-6)

Consequently, the rejections to independent claims 1, 16, 26, and 33 are believed overcome. In light of the above remarks relating to independent claims the remaining dependent claims are believed allowable for at least the reasons noted above. Applicants believe that all pending claims are allowable and respectfully request a Notice of Allowance for this application from the Examiner. Should the Examiner believe that a telephone conference would expedite the prosecution of this application, the undersigned can be reached at the telephone number set out below.

Respectfully submitted,
BEYER WEAVER & THOMAS, LLP

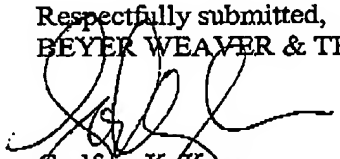

Godfrey K. Kwan
Reg. No. 46,850

P.O. Box 70250
Oakland, CA 94612-0250
(510) 663-1100